

# ***ON POINT***

*High  
Performance  
Computing  
News from  
SSC San Diego*



## **SSC San Diego**

### **High Performance Computing DoD Distributed Center**

SSC San Diego is one of the original sites for the Department of Defense High Performance Computing Modernization Program (DoD HPCMP). Established in 1993 and located at the Point Loma complex in San Diego, the SSC San Diego Distributed Center (DC) has been host to two types of scalable parallel HPC systems: two Intel Paragon XP/S systems, used primarily to support classified and unclassified embedded signal- and image-processing applications research and development (R&D); and the Hewlett-Packard/Convex Exemplar SPP-1600/XA, used primarily to support command and control software development. The DoD HPCMP funded the majority of the costs to establish the SSC San Diego Distributed Center and has provided several increments of funding to upgrade the systems. The most recent upgrade began with DoD acceptance of a 1997 proposal to leapfrog current HPC technologies in conformance with the Navy's Information Technology for the 21st Century (IT-21) and the DoD's 2010 visions. With a Fiscal Year (FY) 98 solicitation and contract award and an additional FY 99 acquisition funding award, the SSC San Diego DC is in the process of replacing its current high performance computing systems with the next generation of supercomputing. These HPC systems will also be scalable and parallel and will be using commodity, commercial off-the-shelf (COTS) hardware and operating system technology to bring this capability to fleet and DoD operational systems of the future.

This upgrade of our DC will bring identical computational environments to our classified and unclassified facilities and will occur in two distinct phases. Phase One will replace the Intel Paragon and Convex Exemplar systems with two Hewlett-Packard V2500 systems and a Windows NT Netserver system. These computational subsystems will be linked to storage subsystems via fibre channel storage area networks. The storage subsystems consist of EMC disk arrays, H-P robotic tape libraries, and hierarchical storage management software that will automatically handle data archival and backup. The current Exemplar system will continue to be available for use until one of the new V2500 systems replaces it; the Paragons will continue to operate as a DoD DC asset through the end of FY 99.

Phase Two of the upgrade will equip the DC with four Hewlett-Packard/Intel IA-64-based systems when those processors become commercially available. These Hewlett-Packard/Intel systems will be capable of running both traditional UNIX and Microsoft Windows NT (WinNT) operating system software and applications. Each of the SSC San Diego HPC systems is connected to campus-wide, internal high-capacity networks. The systems connect directly to both the classified and unclassified asynchronous transfer mode (ATM) optical network backbones that link all of SSC San Diego's major sites and facilities. Together, these computing systems and networks provide SSC San Diego's scientific and engineering researchers with a state-of-the-art HPC environment to support Navy and DoD programs. These programs span the Department of Navy mission areas of command and control, communications, and surveillance, as well as several other technologies and Navy leadership areas.

SSC San Diego is a principal developer of the Defense Research and Engineering Network (DREN). DREN has been a major node on the nationwide high-speed network since its inception. SSC San Diego's role in networking for the DoD provides technical leadership in implementing high bandwidth links between all HPCMP sites and to numerous other DoD science and technology and development, test, and evaluation locations. The DREN node at SSC San Diego provides both secure and unclassified high-capacity connectivity that allows external DoD users to access our DC systems. DREN also permits SSC San Diego's own research, development, test, and evaluation personnel to access other DoD HPC facilities.

The integration of these DC systems and networks in the SSC San Diego HPC environment supports both Navy-specific and Defense-wide mission areas. This integration allows us to join with the other DCs and the Major Shared Resource Centers in contributing to the HPCMP's support of the application of computational science to the technical challenges and systems throughout DoD. The resulting contributions have proven to be very valuable to SSC San Diego, to the Navy's research community, and to the broader DoD Science and Engineering (S&E) and Test and Evaluation (T&E) communities.

HPC systems in the SSC San Diego DC support both classified and unclassified processing. The shared-memory architecture (Hewlett-Packard/Convex Exemplar SPP-1600) is the central component of the Command and Control High Performance Computing Research Facility (CCHPCRF). The distributed-memory architecture (Intel Paragon XP/S-25) is the principal component of the real time, embedded high performance computing facility (RTEHPCF), and processing is supported at both classification levels.

Access to SSC San Diego's DC systems is available to employees and contractors of DoD RDT&E centers and other government agencies under rules specified by the DoD HPC Modernization Program Office. Access requires submission of a completed DoD account application to the appropriate DoD approval authority. The form and additional information may be obtained from the HPCMP World Wide Web site at <http://www.hpcmo.hpc.mil/>. Additional details about the SSC San Diego DC, the CCHPCRF, and the RTEHPCF can be found on our DC Web site at <http://www.sscsd.hpc.mil/dod/>.

SSC San Diego scientists and engineers have been using local HPC capabilities as well as those available at other DoD HPCMO centers to solve complex problems for the Navy and DoD in the primary charter areas of command and control, communications, and surveillance. The following six success stories (detailed in the May 1998 "High Performance Computing Contributions to DoD Mission Success") highlight SSC San Diego's ability to deliver technological advantage to the Warfighter.

## Modeling Contaminant Chemistry in the Environment

J. D. Kubicki and S. E. Apitz

**Research Objective:** To answer basic research questions involved with reducing costs in the management of contaminated sediments. Molecular modeling of organic contaminants in the environment is being used. Knowledge of chemical interactions between organic contaminants and sediments leads to increased understanding of the risk and fate associated with petroleum products released near naval shipyards.

**Results:** To predict the effects of adsorption on chemical stabilities and toxicities of PAHs, three benzene- $[\text{Al}(\text{OH})_{3+x}]^{3-x}$  dimers were energy-minimized with the HF/3-21G\*\* and B3LYP/6-31G\* basis sets. Naphthalene, pyrene, and dibenzo[b,i]anthracene were modeled with an Al-hydroxide surface cluster,  $\text{Al}_9\text{O}_{34}\text{H}_{41}$ , with PM3 semi-empirical quantum mechanical calculations. In some cases, such as the benzene- $[\text{Al}(\text{OH})_2(\text{OH}_2)]^+$  dimer calculations with HF/3-21G\*\* and B3LYP/6-31G\*, the changes in the chemical stability can be on the order of 10%. Thus, significant shifts in molecular stability could occur with adsorption. However, these shifts decrease the chemical stability with regard to degradation. If this is the case for adsorbed benzene and PAHs, then these organic contaminants would be more susceptible to degradation in the adsorbed state, rather than more resistant to degradation, which is what is commonly observed.

**Significance:** Although fairly strong interactions can occur between PAHs and a mineral surface, the chemical bonding is not strong enough to significantly change the chemical stability and toxicity of the contaminants. Furthermore, bonding energetics between organic contaminants and mineral surfaces are an order of magnitude weaker than the bonding between natural organic matter and mineral surfaces. Thus, organic contaminants will not interact directly with mineral surfaces if natural organic matter is already present. These models suggest that decreased biodegradability of organic contaminants is due to a physical sequestration phenomenon within the sediment matrix. Hence, the environmental risk associated with sediments contaminated by PAHs is lower than would be predicted by models. Consequently, cleanup standards for contaminated sediments may be set higher in some instances, which would decrease the volume of sediments requiring remediation and result in significant cost savings.




Image of pyrene, an organic contaminant found in petroleum products, adsorbed onto an aluminum hydroxide mineral surface as modeled with quantum mechanics calculations (blue: carbon; red: hydrogen; green: oxygen; yellow: aluminum). Computations like this increase our understanding of the bioavailability and fate of organic contaminants in marine sediments with the goal of decreasing costs in the management of contaminated sediments.


## High-Speed Radar Imaging for Airborne Target Identification

C. V. Tran and J. R. Evans

**Research Objective:** To make the Noncooperative Airborne Target Identification Code (NATRIC) portable and to obtain a quick turnaround to foster further research and analysis in the radar algorithm-development community. NATRIC is a large suite of computationally intensive software that has been made available by the Office of Naval Research to government researchers and contractors for further studies. It is the cumulative result of more than a decade of research to solve difficult signal and image processing problems in Inverse Synthetic Aperture Radar (ISAR). The suite, however, takes an average of one hour on a Sun SPARCstation 10 to process one second of data.

**Results:** Turnaround in processing time for one second of data was shortened from 60 minutes on a Sun SPARCstation 10 to between 5 and 6 minutes on the Convex Exemplar SPP-1600, depending on the data segment, when using between two to six processors. (This includes the run-time of one component that needed to be processed on a uniprocessor SPARCstation.)

**Significance:** Experiences by U.S. Naval forces in the Persian Gulf War and previously in the Vincennes disaster brought into focus a serious surveillance problem—namely long-range noncooperative target recognition for air threat identification. Radar imaging algorithms can solve this problem; however, they are extremely computationally intensive. This reduction in turnaround computation time from 60 minutes to 5 minutes will allow defense scientists to experiment with various algorithmic nuances and see the results in an interactive mode instead of an overnight batch-processing mode. This should provide a dramatic increase in productivity, which will result in improved algorithms and ultimately in improved air-threat identification systems.



A range/cross-range plot from NATRIC processing with a superimposed outline of the NATRIC-estimated most-probable target.



## Active Sonar Adaptive Beamformer Performance


R. Hiding and J. Lockwood

**Research Objective:** To significantly improve quantification of performance gains of adaptive spatial processing in sonar and radar systems. Adaptive spatial processing has provided significant advancements in interference rejection in sonar and radar systems. However, quantifying performance gains is difficult because performance is scenario-dependent and both signal and interference rejection can occur. A method of performance (MOP) must take into account detection performance as well as interference reduction. Robust, statistical performance measurement is further complicated by the sparse set of target detection opportunities in active surveillance systems and the computationally intensive nature of the algorithms.


**Results:** Robust, statistical performance improvement over a large data set was demonstrated for the active adaptive beamformer (ABF). The evaluation was done in a system context providing direct, meaningful comparisons for active surveillance sonar system designers. The use of the scalable programming environment, a scalable signal processing application programming interface, allowed throughputs twice as fast as real time for ABF and seven times faster than real time for CBF, and significantly reduced development time.

**Significance:** The application of ABF to active sonar surveillance will improve detection and classification of difficult threats in highly cluttered littoral warfare scenarios.

This work was done jointly with T. Baarnard and J. Smigel from Lockheed-Martin, Syracuse, NY.



Typical receiver operating characteristic curve comparing ABF performance with that of CBF.




Average false alarm rate reduction of adaptive vs conventional beamforming for two adaptive processing parameters.

## Bistatic Target Strength Prediction from Limited Data

A. K. Kevorkian, D. Barach and G. W. Benthien

**Research Objective:** To develop a fast parallel algorithm that enables the extension of the frequency range relevant to bistatic target strength prediction to regions of greater interest to the Navy. Knowledge of bistatic target strength is of increasing importance in sonar systems. Full-scale measurements of monostatic target strength are expensive and difficult, and become impractical for general bistatic geometries. A numerical model developed by Schenk, Benthien, and Barach uses measured monostatic and limited bistatic data to estimate the surface field and propagate it to the farfield for full bistatic geometries. Existing data sets cannot be fully exploited using present solution techniques.

**Results:** The most compute-intensive component of the bistatic target strength model involves the least-squares solution of large, sparse, rank-deficient, overdetermined matrices. Available conventional algorithms, such as those in the linear algebra package LAPACK, deal with the rank deficiency but take no advantage of the sparsity in these models. Therefore, they are time consuming, and this makes the method impractical for existing data sets. We developed a new complete orthogonal factorization that takes full advantage of the structure and numerical properties in sparse, rank deficient, overdetermined systems of the type arising in bistatic target strength prediction models. By applying this decomposition method to a range of highly ill-conditioned rank-deficient bistatic target strength prediction models, we have obtained a five- to seven-fold improvement over the compute time required by the linear algebra package LAPACK without exploiting any parallelism. By exploiting parallelism in one key



Comparison of estimated and true monostatic and bistatic target strength surfaces.

part of our method, we have obtained improvements over LAPACK from nine- to fifteen-fold. Further speedups can be realized by exploiting parallelism in all parts of the new method.

**Significance:** Since the size of the problem increases rapidly with frequency, the speedups gained from the new method will enable the extension of the frequency range in bistatic target strength prediction models to regions of greater interest to the Navy.

## Scalable Prototyping of Embedded Signal Processing Systems


P. Partow and D. M. Cottel

**Research Objective:** To develop a portable version of the SSC San Diego Scalable Programming Environment (SPE) and to provide and support versions for various HPC systems and for networks of workstations. DoD embedded signal processing applications are currently developed as conventional sequential programs and then rewritten as parallel programs for a specific embedded architecture. The SPE provides a way to develop scalable, parallel signal and image processing applications that can, by simple recompilation, be ported to various HPC architectures and ultimately to embedded systems. The concepts have been proven on the Intel Paragon where the SPE is in production use.

**Results:** Under the DoD HPC Modernization Program, with very little change to the user-visible programming interface, the SPE was redesigned and reimplemented to use MPI-1 for communication between processes.

Initial beta test versions of the SPE have been distributed to potential users of Windows NT, Sun Solaris, and Silicon Graphics IRIX systems. Projects involving Synthetic Aperture Radar (SAR) at SSC San Diego, Automatic Target Recognition (ATR) at Aeronautical Systems Center at Wright Patterson Air Force Base, and Image Compression and ATR at the Army Night Vision Laboratory are developing SPE-based scalable, parallel modules that will be interconnected in a planned tri-service interoperability demonstration. SSC San Diego and Sanders are collaborating to develop an implementation of the SPE for the Sanders Embedded High Performance Scalable Computing System. Khoros Research, Inc., the developer of Khoros, is considering ways of incorporating SPE technology into the Advanced Khoros project currently funded by the Defense Advanced Research Projects Agency (DARPA).

**Significance:** With availability on a wide variety of platforms, the SPE has the potential to be the foundation on which parallel, scalable software modules can be provided to the signal and image processing community. Because of the SPE's support for modular programming, the various modules can be freely interconnected in much the same style as Khoros. The resulting applications can be directly ported from workstations, to HPC systems, and to any embedded system for which the SPE has been implemented.



Functional and data-flow diagram of an operational active sonar processor implemented as a parallel, scalable HPC program. The application will run on an arbitrary number of nodes of the Intel Paragon.


## Parallelization of NEC-BSC for Antenna Performance Prediction for Improved C4I

Dr. C. W. Manry Jr., J. Strauch, Dr. J. H. Meloling, and A. M. Lu


**Research Objective:** To determine if Ohio State University's Numerical Electromagnetic Code – Basic Scattering Code (NEC-BSC) could be parallelized by using SSC San Diego's Intel Paragon high performance computer. The NEC-BSC is based on the Uniform Geometrical Theory of Diffraction (UTD), which allows structures to be modeled as a set of geometric entities. The number of geometric entities on a large ship (LHD 7 class), coupled with the large number of antenna sources required to perform a calculation, is a computational challenge.

**Results:** A computation that took 13 hours on a single processor SGI Indigo-2 now only takes 1 hour with the Intel Paragon using 181 nodes. For problems involving more sources, the speed enhancement appears to be greater.

**Significance:** Due to this speed enhancement, results can be examined, changes made, and new tests performed in less than one-tenth the time. The increased efficiency of the parallelized codes on the Intel Paragon has made the analysis and design efforts at SSC San Diego more thorough. Antenna performance is therefore improved, thus allowing for increased range and performance of C4I systems for the Warfighter. Also, since the NEC-BSC code was parallelized by using Intel NX (which ports to modern MPI machines such as the Cray T3E, SGI Origin 2000, and the IBM SP), the benefit will continue as the Intel Paragon is phased out of DoD HPC centers.



NEC-BSC volumetric pattern of the JTIDS antenna (AS-4400/URC) on the LHD-1 – stern view.



NEC-BSC ray visualization between relocated antennas 3-5 and 3-6 (AS-2810/SRCs) on the CG 47.

## ***SSC San Diego Distributed Center***

SSC San Diego is one of 13 Distributed Centers funded by the DoD Modernization Program. The SSC San Diego Distributed Center (DC) hosts two types of scalable parallel HPC systems: two Intel Paragon XP/S systems and a Hewlett-Packard/Convex Exemplar SPP-1600/XA. High-speed connectivity to these local computers and all other HPC computers linked to Defense Research and Engineering Network (DREN) and other national networks is made possible through the SSC San Diego networking infrastructure.

SSC San Diego DC is in the process of replacing its current high performance computing systems with the next-generation of supercomputing. This upgrade will bring identical computational environments to our classified and unclassified facilities and is occurring in two distinct phases. Phase One will replace the Intel Paragon and Convex Exemplar systems with two Hewlett-Packard V2500 systems and a Windows NT Netserver system. Phase Two of the upgrade will equip the center with four Hewlett-Packard/Intel IA-64 based systems when those processors become commercially available.



SSC San Diego  
San Diego, CA 92152-5001

Reviewed and approved by

Executive Officer/  
Base Operations Manager  
SSC San Diego

SD 128, Rev. 1  
April 1999

Approved for public release;  
distribution is unlimited.